

Serial No. 09/996,074

Page 8 of 17

REMARKS

This response is intended as a full and complete response to the non-final Office Action mailed on February 22, 2005. In the Office Action, the Examiner notes that claims 1-18 and 20-30 are pending and rejected. By this response, claims 1 and 17 have been amended, and claims 2-16, 18 and 20-30 continue unamended.

In view of the foregoing amendments and the following discussion, the Applicants submit that none of the claims now pending in the application are anticipated or obvious under the provisions of 35 U.S.C. §§102 or 103. Thus, the Applicants believe that all these claims are now in allowable form.

It is to be understood that the Applicants, by amending the claims, do not acquiesce to the Examiner's characterizations of the art of record or to the Applicants' subject matter recited in the pending claims. Further, the Applicants are not acquiescing to the Examiner's statements as to the applicability of the prior art of record to the pending claims by filing the instant responsive amendments.

REJECTIONS

35 U.S.C. §102

Claims 1-4, 10 and 11

The Examiner has rejected claims 1, 2, 3, 4, 10 and 11 under 35 U.S.C. §102(b) as being anticipated by Bailey (USPN 5623595A, hereinafter "Bailey"). The Applicants respectfully traverse the rejection.

The Applicants' independent claim 1 recites:

"A method for streaming content striped in RAID 5 format from an array of disk drives to a plurality of subscribers to minimize disruptive service from a disk drive failure, said method comprising:
accessing content data striped in said RAID 5 format, on an extent-by-extent basis, from a plurality of disk drives configured in an array without reading the parity extents;
streaming the content data to the plurality of subscribers on an extent-by-extent basis, sequentially, from the plurality of disk drives;
detecting an actual disk drive failure; and
transitioning to a stream regeneration mode of operation
comprising:

344171-1

Serial No. 09/996,074

Page 9 of 17

reading the content data contemporaneously from all extents in a parity group associated with a failed disk drive;
regenerating a failed portion of the content data from a failed extent in the parity group corresponding to the failed disk drive; and
streaming the content data in the parity group to the plurality of subscribers, extent-by-extent, immediately following the regenerating of the content data from the failed extent in the parity group." (emphasis added).

Anticipation requires the presence in a single prior art reference disclosure of each and every element of the claimed invention, arranged as in the claim" (*Lindemann Maschinenfabrik GmbH v. American Hoist & Derrick Co.*, 730 F.2d 1452, 221 USPQ 481, 485 (Fed. Cir. 1984)(citing *Connell v. Sears, Roebuck & Co.*, 722 F.2d 1542, 220 USPQ 193 (Fed. Cir. 1983).

The Bailey reference discloses that:

[D]ata is stored and retrieved from a redundancy group of data storage units in increments of a stripe at a time, which in this example totals 256 KB, stored in four blocks of 64 KB each. The fifth block of each stripe is a parity block that contains the XOR'd sum of the four data blocks in the stripe. To read the portion of the video program stored in stripes "A" to "F" of FIG. 2, read requests are first simultaneously issued for all the blocks in stripe "A", namely blocks A1, A2, A3, A4 and A5. These read requests, depending on the embodiment, may be issued by a process running in a data processing device or by an intelligent storage unit controller."

A read request for block A1 is issued to storage unit SU1, a read request for block A2 is issued to storage unit SU2, a read request for block A3 is issued to storage unit SU3, a read request for block A4 is issued to storage unit SU4, and a read request for block A5 is issued to storage unit SU5, all simultaneously. The requested data blocks are not always delivered at the same time. Instead, they are spread over some time period. In one embodiment of the present invention, as soon as a block is received, it is stored in a buffer, and a check is made to determine whether received block is the next to last block in the stripe (i.e. the fourth block in this example). If the most recently received block is the next to last block (indicating that all but one of the blocks of the stripe have been received), a further check is made to determine whether all of the received blocks are data blocks or whether, instead, one of them is the parity block. If the first four blocks are all data blocks, all data blocks of the stripe have been received, and they are delivered to the requesting device or process. If one of the first four blocks is the parity block for the stripe, then the block that is still outstanding is a data block. Rather than continue to wait for the

344171-1

Serial No. 09/996,074

Page 10 of 17

outstanding block, the present invention reconstructs the missing block by XOR'ing the three data blocks and the parity block that have already been received. The three received data blocks, together with the XOR'd sum of those three blocks and the parity block, are delivered to the requesting device or process. Because a stripe contains four data blocks and a parity block that contains the XOR'd sum of the four data blocks, the XOR'd sum of any four out of five blocks of a stripe is the same as the fifth block, regardless of the order in which the blocks are received. Accordingly, with the present invention, there is no need to wait for the fifth and last block to be delivered, regardless of the reason for the fifth block's delay. (See Bailey, column 8, line 61 to column 9, line 39).

The Bailey reference differs from the Applicants' invention since the Bailey reference fails to teach, or even suggest "accessing content data striped in said RAID 5 format, on an extent-by-extent basis, from a plurality of disk drives configured in an array without reading the parity extents." Bailey retrieves all the data blocks and parity block a stripe at a time. After retrieving 4 of the 5 blocks, the final block is generated (if necessary) and the data is transmitted regardless of the reason for the delay.

On the other hand, the Applicants' invention does not even read the parity information. "Only the data extents are read in a parity group 404, while the parity extent 408 containing the parity information and spare extent 410 are not read." (See Specification, page 11, lines 31-33). During normal mode of operation, it is unnecessary to spend valuable resources retrieving the parity information from the disk drive array. The Applicants' invention achieves this bandwidth efficiency by only accessing content data. However, Bailey is directed toward improving the speed of the system so it always accesses the parity information in order to generate the last data block when necessary. As a result, Bailey does not disclose "accessing... without reading the parity extents." Therefore, the Bailey reference fails to teach each and every element of the claimed invention, as arranged in the claim.

As such, the Applicants submit that independent claim 1 is not anticipated under 35 U.S.C. §102 and is fully patentable thereunder. Furthermore, claims 2, 3, 4, 10 and 11 depend, either directly or indirectly, from independent claim 1 and recite additional limitations thereof. As such and for at least the same reasons, the Applicants submit that these dependent claims are not anticipated under 35 U.S.C. §102 and are fully

344171-1

Serial No. 09/996,074

Page 11 of 17

patentable thereunder. Therefore, the Applicants respectfully request that the rejections of claims 1, 2, 3, 4, 10 and 11 be withdrawn.

35 U.S.C. §103

Claims 5-9, 12-16

The Examiner has rejected claims 5-9 and 12-16 under 35 U.S.C. §103(a) as being unpatentable over Bailey in view of Peters et al. (USPN 6415373, hereinafter "Peters"). The Applicants respectfully traverse the rejection.

The test under 35 U.S.C. §103 is not whether an improvement or a use set forth in a patent would have been obvious or non-obvious; rather the test is whether the claimed invention, considered as a whole, would have been obvious. Jones v. Hardy, 110 USPQ 1021, 1024 (Fed. Cir. 1984) (emphasis added). Thus, it is impermissible to focus either on the "gist" or "core" of the invention, Bausch & Lomb, Inc. v. Barnes-Hind/Hydrocurve, Inc., 230 USPQ 416, 420 (Fed. Cir. 1986). Moreover, the invention as a whole is not restricted to the specific subject matter claimed, but also embraces its properties and the problem it solves. In re Wright, 6 USPQ 2d 1959, 1961 (Fed. Cir. 1988) (emphasis added).

Claims 5-9 and 12-16 are dependent directly or indirectly upon independent claim 1 and include additional features thereof.

Claim 5 (and similarly claims 6-9 and 12-16), which included the features of independent claim 1, recites in part:

"A method for streaming content striped in RAID 5 format from an array of disk drives to a plurality of subscribers to minimize disruptive service from a disk drive failure, said method comprising:
accessing content data striped in said RAID 5 format, on an extent-by-extent basis, from a plurality of disk drives configured in an array without reading the parity extents;
streaming the content data to the plurality of subscribers on an extent-by-extent basis, sequentially, from the plurality of disk drives;
detecting an actual disk drive failure; and
transitioning to a stream regeneration mode of operation
comprising:

344171-1

Serial No. 09/996,074

Page 12 of 17

reading the content data contemporaneously from all extents in a parity group associated with a failed disk drive;
regenerating a failed portion of the content data from a failed extent in the parity group corresponding to the failed disk drive; and
streaming the content data in the parity group to the plurality of subscribers, extent-by-extent, immediately following the regenerating of the content data from the failed extent in the parity group."

The Bailey reference does not teach or suggest: "accessing content data striped in said RAID 5 format, on an extent-by-extent basis, from a plurality of disk drives configured in an array without reading the parity extents." As discussed above, with respect to the 35 U.S.C. §102 rejection, the Bailey reference merely discloses retrieves all the data blocks and parity block a stripe at a time. After retrieving 4 of the 5 blocks, the final block is generated (if necessary) and the data is transmitted regardless of the reason for the delay.

Furthermore, the Peters reference fails to bridge the substantial gap as between the Bailey reference and the Applicants' invention. Specifically, the Peters reference discloses "[a]fter a storage unit fails, a new storage unit may be installed in its place, with lost data restored, or the lost data may be recreated and distributed over the remaining storage units." (See Peters, column 15, lines 37-40). It does not disclose, teach or even suggest: "accessing content data striped in said RAID 5 format, on an extent-by-extent basis, from a plurality of disk drives configured in an array without reading the parity extents."

As stated in the MPEP 2143.03,

To establish prima facie obviousness of a claimed invention, all the claim limitations must be taught or suggested by the prior art. In re Royka, 490 F.2d 981, 180 USPQ 580 (CCPA 1974). "All words in a claim must be considered in judging the patentability of that claim against the prior art." In re Wilson, 424 F.2d 1382, 1385, 165 USPQ 494, 496 (CCPA 1970). If an independent claim is nonobvious under 35 U.S.C. 103, then any claim depending therefrom is nonobvious. In re Fine, 837 F.2d 1071, 5 USPQ2d 1596 (Fed. Cir. 1988).

As stated above, neither of the references discloses, teaches or suggests the limitation of "accessing content data striped in said RAID 5 format ... without reading parity extents." Therefore, even if the two references could somehow be forcibly combined,

344171-1

Serial No. 09/996,074

Page 13 of 17

the combination would still lack "accessing content data striped in said RAID 5 format ... without reading parity extents." Accordingly, the Applicants' invention is not obvious in view of the cited references. Therefore, the combined references fail to teach or suggest Applicants' invention as a whole.

As such, the Applicants submit that dependent claims 5-9 and 12-16 are not obvious and fully satisfies the requirements of 35 U.S.C. §103 and is patentable over Bailey in view of Peters. Therefore, the Applicants respectfully request that the Examiner's rejection of claims 5-9 and 12-16 be withdrawn.

Claims 17-22

The Examiner has rejected claims 17-22 under 35 U.S.C. §103(a) as being unpatentable over Bailey in view of Brady et al. (USPN 5727144, hereinafter "Brady"). The Applicants respectfully traverse the rejection.

The Applicants' independent claim 17 recites:

"A method for streaming content striped in RAID 5 format from an array of disk drives to a plurality of subscribers to minimize disruptive service from a disk drive failure, said method comprising:

- accessing content data striped in said RAID 5 format, on an extent-by-extent basis, from a plurality of disk drives configured in an array without reading the parity extents;
- streaming the content data to the plurality of subscribers on an extent-by-extent basis, sequentially, from the plurality of disk drives;
- predicting a disk drive failure;
- writing content data from a the disk drive predicted to fail to spare extents on non-failed disk drives in the array;
- detecting at least one of an actual failure and removal of the disk drive predicted to fail; and
- transitioning, in response to the detecting step, a recovery-carousel-serving mode of operation comprising:
 - streaming psuedo-sequentially, extent-by-extent, content data of each parity group to the plurality of subscribers, where the regenerated content data in a spare extent of each parity group is streamed out of sequence;
 - wherein in an instance where the disk drive predicted to fail fails prior to said writing step, said method further comprises
 - transitioning to a stream regeneration mode of operation comprising:

344171-1

Serial No. 09/996,074

Page 14 of 17

reading the content data contemporaneously from all extents in a parity group;
regenerating a failed portion of the content data from a failed extent in the parity group corresponding to the failed disk drive; and
streaming the content data in the parity group to the plurality of subscribers, extent-by-extent, immediately following the regenerating of the content data from the failed extent in the parity group." (emphasis added).

As discussed above, the Bailey reference fails to disclose each and every element of the claimed invention, as arranged in the claim. In particular, the Bailey reference does not teach or suggest: "accessing content data striped in said RAID 5 format, on an extent-by-extent basis, from a plurality of disk drives configured in an array without reading the parity extents." The Bailey reference merely discloses retrieves all the data blocks and parity block a stripe at a time. After retrieving 4 of the 5 blocks, the final block is generated (if necessary) and the data is transmitted regardless of the reason for the delay.

Furthermore, the Brady reference fails to disclose, teach or suggest accessing content data striped in said RAID 5 format, on an extent-by extent bases, from a plurality of disk drives configured in an array without reading the parity extents. Specifically, the Brady reference discloses:

"[I]n a data processing system employing a disk array, prediction of a possible failure of a disk drive initiates copying of the data away from the potentially failing disk drive to a spare disk drive before the failing drive actually fails. If the disk drive does fail before the copying of the contents to a spare disk drive is completed, rebuilding of the remaining contents within the failing disk drive is performed. When a warning of an imminent disk failure is received, the present invention proceeds to copyaway state 21, wherein data is copied away from failing drive 110 to spare drive 112. When copyaway completes, the present invention returns to normal mode 20. It is well-known by those skilled in the art how to copy contents of drive 110 to drive 112.

Failing disk 110 may fail before the copyaway procedure is completed. In that case, a transition to rebuild mode 24 is performed, wherein a rebuilding of those portions of disk drive 110 that had not yet been copied away is performed.

Another error situation shown is when a disk other than the one for which the warning was initiated fails without warning (e.g., disk 111). For example, assume that a warning was initiated because of a pending

344171-1

Serial No. 09/996,074

Page 15 of 17

failure of disk 110, which initiated I/O controller 180 to begin copying away from disk 110 to disk 112 in copyaway mode 21. During copyaway of disk 110, suppose that disk 111 fails. In this instance, a transition is made to the high priority copyaway state 22, where copyaway is given high priority and is no longer done in the background. This is because if disk 110 were to fail before copyaway of 110 completes, data would be lost. If copyaway does complete before disk 110 fails, a transition is made to rebuild state 24 where the contents of failed disk 111 are rebuilt." (See, Brady, Abstract, column 4, lines 15-38.)

It does not disclose, teach or even suggest: "accessing content data striped in said RAID 5 format, on an extent-by-extent basis, from a plurality of disk drives configured in an array without reading the parity extents." As stated above, neither of the references discloses, teaches or suggests the limitation of "accessing content data striped in said RAID 5 format, on an extent-by-extent basis, from a plurality of disk drives configured in an array without reading the parity extents." Therefore, even if the two references could somehow be operably combined, the combination would still lack "accessing content data striped in said RAID 5 format, on an extent-by-extent basis, from a plurality of disk drives configured in an array without reading the parity extents." Accordingly, the Applicants' invention is not obvious in view of the cited references. Therefore, the combined references fail to teach or suggest Applicants' invention as a whole.

As such, the Applicants submit that claim 17 is not obvious and fully satisfies the requirements of 35 U.S.C. §103 and is patentable thereunder. Claims 18-22 depend, directly or indirectly, from claim 17 and recite similar features thereof. As such, and for at least the same reasons as discussed above, the Applicants submit that claims 18-22 are also not obvious and fully satisfy the requirements of 35 U.S.C. §103 and are patentable thereunder. Therefore, the Applicants respectfully request that the rejections be withdrawn.

Claims 23-27

The Examiner has rejected claims 23-27 under 35 U.S.C. §103(a) as being unpatentable over Bailey in view of Brady in further view of Peters.

344171-1

Serial No. 09/996,074
Page 16 of 17

For at least the reasons discussed above, Applicants submit that independent claim 17 is not obvious and is patentable over Bailey in view of Brady. Claims 23-27 depend, directly or indirectly, upon independent claim 17 and recite similar features thereof. As such, and for at least the same reasons as discussed above, the Applicants submit that claims 23-27 are also not obvious and fully satisfy the requirements of 35 U.S.C. §103 and are patentable over Bailey in view of Brady.

Furthermore, the Peters reference fails to bridge the substantial gap as between Bailey and Brady and the Applicants' invention. Specifically, as discussed above, the Peters reference discloses installing a new disk drive to replace a failed disk drive in an array (See Peters, column 15, lines 37-40). It does not disclose, teach or suggest accessing content data without reading parity extents.

Even if the three references could somehow be operably combined, the combination would still lack the limitation of "accessing content data striped in said RAID 5 format, on an extent-by-extent basis, from a plurality of disk drives configured in an array without reading the parity extents." Therefore, the combined references fail to teach or suggest the Applicants' invention as a whole.

As such, the Applicants submit that independent claim 23 is not obvious and fully satisfies the requirements of 35 U.S.C. §103 and is patentable over Bailey and Brady in view of Peters. For at least the same reasons as discussed above, the Applicants submit that dependent claims 24-27 are also not obvious and fully satisfy the requirements of 35 U.S.C. §103 and are patentable thereunder. Therefore, the Applicants respectfully request that the Examiner's rejection of claims 23-27 be withdrawn.

SECONDARY REFERENCES

The secondary references made of record are noted. However, it is believed that the secondary references are no more pertinent to the Applicants' disclosure than the primary references cited in the Office Action. Therefore, the Applicants believe that a detailed discussion of the secondary references is not necessary for a full and complete response to this office action.

344171-1

Serial No. 09/996,074
Page 17 of 17

Conclusion

Thus, the Applicants submit that claims 1-18 and 20-30 are in condition for allowance. Accordingly, both reconsideration of this application and its swift passage to issue are earnestly solicited.

If, however, the Examiner believes that there are any unresolved issues requiring adverse final action in any of the claims now pending in the application, it is requested that the Examiner telephone Eamon J. Wall, Esq. at (732) 530-9404 so appropriate arrangements can be made for resolving such issues as expeditiously as possible.

Respectfully submitted,

5/23/05

E J Wall
EAMON J. WALL, Attorney
Reg. No. 39,414
(732) 530-9404

Moser, Patterson & Sheridan, LLP
Attorneys at Law
595 Shrewsbury Avenue, Suite 100
Shrewsbury, New Jersey 07702

344171-1